

In the Claims:

Following is a complete listing of the claims pending in the application, as amended:

1. (Original) A method in a computing system for processing a relational database query, comprising;
receiving the relational database query;
constructing a multidimensional database query based on the received relational database query; and
submitting the constructed multidimensional database query for execution against a multidimensional data source.
2. (Original) The method of claim 1, further comprising:
receiving, in response to submitting the multidimensional database query, a multidimensional database query result; and
constructing a relational database query result based on the received multidimensional database query result.
3. (Original) A method in a computing system for processing a relational database query, comprising;
receiving the relational database query, the received relational database query being drawn against a relational model of a multidimensional data source;
using a relational-to-multidimensional mapping contained by the model together with relational/multidimensional equivalency logic to construct a multidimensional database query based on the received relational database query; and
submitting the constructed multidimensional database query for execution against the modeled multidimensional data source.

4. (Original) The method of claim 3 wherein the multidimensional query is constructed in MDX.

5. (Original) The method of claim 4 wherein the relational query is expressed in SQL.

6. (Original) The method of claim 3 wherein the relational query specifies a detail filter against the relational model, and wherein the constructed multidimensional query specifies that the detail filter be applied to the modeled multidimensional data source.

7. (Original) The method of claim 6 wherein the relational query is expressed in SQL, and wherein the detail filter specified by the relational query is an SQL WHERE clause.

8. (Original) The method of claim 3 wherein the relational query specifies an aggregation function against the relational model, and wherein the constructed multidimensional query specifies that the aggregation function be applied to the modeled multidimensional data source.

9. (Original) The method of claim 8 wherein the relational query is expressed in SQL, and wherein the aggregation function specified by the relational query is an SQL GROUP BY clause.

10. (Original) The method of claim 3 wherein the relational query specifies a summary filter against the relational model, and wherein the constructed multidimensional query specifies that the summary filter be applied to the modeled multidimensional data source.

11. (Original) The method of claim 10 wherein the relational query is expressed in SQL, and wherein the summary filter specified by the relational query is an SQL HAVING clause.

12. (Original) The method of claim 3 wherein the relational query specifies a detail filter against the relational model having selected predicates, and wherein the constructed multidimensional query specifies, for each of the selected predicates that can be applied against the modeled multidimensional data source before a crossjoin operation is performed, applying the selected predicate against the modeled multidimensional data source as early as possible.

13. (Original) The method of claim 3 wherein the relational query specifies a detail filter against the relational model having selected predicates, and wherein the constructed multidimensional query specifies, for each of the selected predicates that can be applied against the modeled multidimensional data source before a crossjoin operation is performed, applying the selected predicate against the modeled multidimensional data source before a crossjoin operation is performed.

14. (Original) The method of claim 3 wherein the relational query specifies performing a selected aggregation function on a selected column of a virtual relational table, the virtual relational table corresponding to a multidimensional data source, the selected column corresponding to a selected measure of the multidimensional data source, the method further comprising:

retrieving metadata identifying an aggregation function used for the selected measure of the multidimensional data source;

determining whether the aggregation function identified by the metadata matches the selected aggregation function; and

if the aggregation function identified by the metadata matches the selected aggregation function, generating a multidimensional query against the

multidimensional data source that relies on the aggregation function performed in the multidimensional data source.

15. (Original) The method of claim 3, further comprising:
receiving, in response to submitting the multidimensional database query, a multidimensional database query result; and
using a relational-to-multidimensional mapping contained by the model together with relational/multidimensional equivalency logic to construct a relational database query result based on the received multidimensional database query result.

16. (Original) The method of claim 3, further comprising:
determining that the received relational database query is drawn against both the relational model of the multidimensional data source and one or more native relational tables; and
constructing a native relational database query based on aspects of the received relational database query drawn against conventional relational tables;
and
submitting the native relational database query for execution against the conventional relational tables,

and wherein the constructed multidimensional database query is based on aspects of the received relational database query drawn against the relational model of the multidimensional data source, the method further comprising:

receiving, in response to submitting the native relational database query, a native relational database query result; and
combining the constructed relational database query result with the received native relational database query result in accordance with the received relational database query.

17. (Original) The method of claim 3, further comprising making information about the model available for use in building the received relational database query.

18. (Original) The method of claim 3, further comprising:
determining that the received relational database query is drawn against both the relational model of the multidimensional data source and one or more native relational tables; and
constructing a native relational database query based on aspects of the received relational database query drawn against conventional relational tables;
and
submitting the native relational database query for execution against the conventional relational tables,

and wherein the constructed multidimensional database query is based on aspects of the received relational database query drawn against the relational model of the multidimensional data source.

19. (Currently Amended) A computer-readable medium whose contents cause a computing system ~~to~~for processing a relational database query by;
receiving the relational database query, the received relational database query being drawn against a relational model of a multidimensional data source;
using a relational-to-multidimensional mapping contained by the model to translate the received relational database query into a multidimensional database query; and
submitting the multidimensional database query for execution against the modeled multidimensional data source.

20. (Currently Amended) The computer-readable medium of claim 19, wherein the contents of the computer-readable medium further cause the computing system to further comprising:

receiving, in response to submitting the multidimensional database query, a multidimensional database query result; and
using a relational-to-multidimensional mapping contained by the model to translate the received multidimensional database query result into a relational database query result.

21. (Original) A computing system for processing a relational database query, comprising;

a query reception subsystem that receives the relational database query, the received relational database query being drawn against a relational model of a multidimensional data source;
a multidimensional query construction subsystem that uses a relational-to-multidimensional mapping contained by the model to construct a multidimensional database query based on the received relational database query; and
a query submission subsystem that submits the constructed multidimensional database query for execution against the modeled multidimensional data source.

22. (Original) The computing system of claim 21, further comprising:

a query result reception subsystem that receives, in response to submitting the multidimensional database query, a multidimensional database query result; and
a relational query result construction subsystem that uses a relational-to-multidimensional mapping contained by the model to construct a relational database query result based on the received multidimensional database query result.

23. (Original) One or more computer memories collectively containing a data source modeling data structure for use in modeling a multidimensional data source in a relational database environment, the data structure comprising:

a schema for one or more virtual relational tables each representing contents of the multidimensional data source; and

one or more mappings between schema components and contents of the multidimensional data source to which they correspond,

such that the schemas contained by the data structure may be used to formulate relational queries against the virtual relational tables,

and such that the schemas and mappings may be used to translate relational queries against the virtual relational tables into multidimensional queries against the multidimensional data source and to translate multidimensional query results from the multidimensional data source into relational query results from the virtual relational tables.

24. (Original) The computer memories of claim 23 wherein, in the multidimensional data source, a particular multidimensional level name occurs in both a first hierarchy of the multidimensional data source and a second hierarchy of the multidimensional data source that is distinct from the first hierarchy,

and wherein the schema for the virtual relational tables specify both a first column corresponding to the occurrence of the multidimensional level name in the first hierarchy and a second column corresponding to the occurrence of the multidimensional level name in the second hierarchy, the metadata for the first column specifying an external name that is the same as the multidimensional level name as well as an internal name, and the metadata for the second column specifying an external name that is the same as the multidimensional level name and an internal name that is distinct from the internal name specified for the first column.

25. (Original) The computer memories of claim 23 wherein the data structure further comprises, for each of a plurality of members of the multidimensional data

source, metadata identifying an aggregation rule applied to the measure in a multidimensional database environment in which the modeled multidimensional data source resides.

26. (Original) One or more computer memories collectively containing a database-type transparency data structure for use in modeling a plurality of multidimensional data source in a relational database environment, the data structure comprising:

for each of the multidimensional data sources, individual source information comprising:

information defining one or more corresponding virtual relational tables,
and

information mapping between components of the virtual relational tables
and contents of the multidimensional data source; and

a single body of relational/multidimensional equivalency logic that may be used to translate a relational query against one or more of the virtual relational tables defined the individual source information for selected multidimensional data sources into a multidimensional query against the selected multidimensional data sources with reference to the individual source information for the selected multidimensional data sources.

27. (Original) One or more computer memories collectively containing a data structure constituting a multidimensional database query against a multidimensional data source, the data structure comprising contents generated by analyzing a relational database query issued against virtual relational tables that model the multidimensional data source,

such that the contents of the data structure may be used to execute the multidimensional database query.

28. (Original) One or more generated data signals collectively conveying a data structure constituting a multidimensional database query against a multidimensional data source, the data structure comprising contents generated by analyzing a relational database query issued against a relational model of the multidimensional data source,

such that the contents of the data structure may be used to execute the multidimensional database query.

29. (Original) One or more computer memories collectively containing a relational database query result data structure, the data structure comprising contents relating to relational tables that model a multidimensional data source, and generated by analyzing a multidimensional database query result produced from the multidimensional data source.

30. (Original) One or more generated data signals collectively conveying a relational database query result data structure, the data structure comprising contents relating to a relational model of a multidimensional data source, and generated by analyzing a multidimensional database query result produced from the multidimensional data source.

31. (Original) A method in a computing system for processing a relational database query, comprising:

receiving the relational database query, the received relational database query being drawn against both a relational model of a multidimensional data source and a native relational table;

converting the received relational database query into (1) a native relational database query against only the native relational table, and (2) a multidimensional database query against the multidimensional data source;

submitting the a native relational database query against the native relational table;

submitting the multidimensional database query against the multidimensional data source; and

combining contents of a first search result produced in response to the native relational database query and a second search result produced in response to the multidimensional database query into a third search result responsive to the received relational database query.

32. (Original) A method in a computing system for generating a multidimensional query, comprising:

receiving a relational query that specifies performing a selected aggregation function on a selected column of a virtual relational table, the virtual relational table corresponding to a multidimensional cube, the selected column corresponding to a selected measure of the multidimensional cube;

retrieving metadata identifying an aggregation function used for the selected measure of the multidimensional cube;

determining whether the aggregation function identified by the metadata matches the selected aggregation function; and

if the aggregation function identified by the metadata matches the selected aggregation function, generating a multidimensional query against the multidimensional cube that relies on the aggregation function performed in the multidimensional cube.

33. (Original) The method of claim 32 wherein the multidimensional query is generated only if the relational query does not specify a filter referencing levels of the multidimensional cube that are below a level of aggregation specified by the relational query.

34. (Original) The method of claim 32 wherein the generated query does not specify performing a crossjoin in order to perform the selected aggregation function.